Current Status of CAS Integration

Brian Yanny

SDSS CAS Review
Fermilab
March 1-2, 2004

Internal SDSS operations support

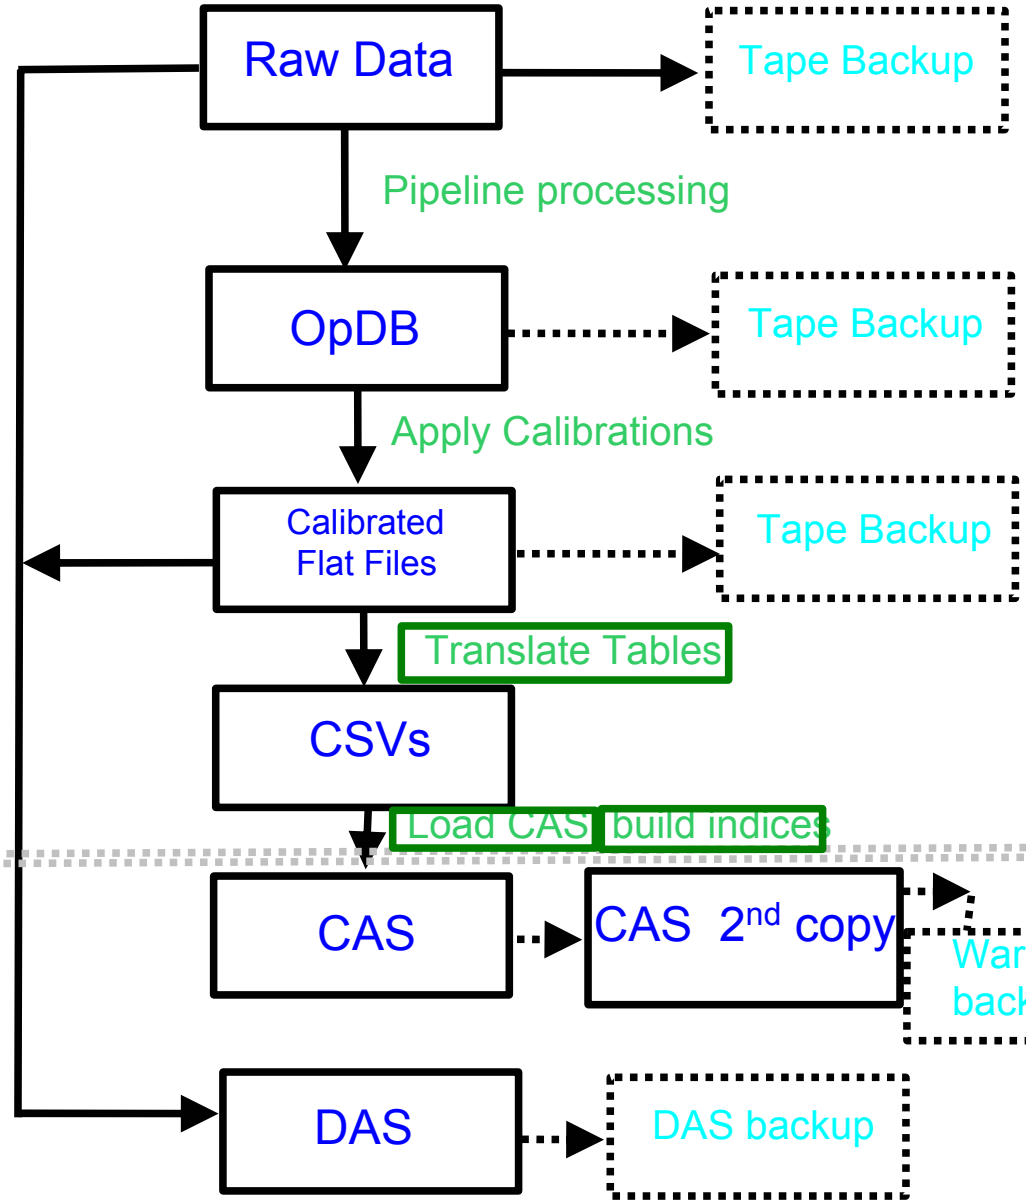

	DR1 2003	DR2 2004	DR5 (full survey) 2006+
	5TB	10TB	20TB (on tape)
	1TB	2TB	5TB per copy
	1TB	2TB	5TB per copy
	(1TB	2TB	5TB – ephemeral)
	1TB	2TB	5TB per copy
	[3 TB	6TB	15TB per copy]

External SDSS Data Distribution Support

Current Status of CAS Integration

1. Current Status of integration into operations environment

Hand-over of Data from 'internal SDSS operations' to 'CAS based data distribution':

Survey operations produce 'calibrated imaging object catalogs',
'extracted spectra with measured parameters'
and a variety of supporting bookkeeping files in FLAT FILE FORMAT.
The format is FITS (astronomers standard self-documenting binary table)
format: 'Flexible Image Transport System', as well as ASCII 'parameter ' files:
(lists of keyword=value and simple tables).

The extent and contents of these files are defined and tracked by
'the data model' which is a set of text web pages which says what's
in each file, with units, and format.

Translate tables: In order to be loaded into the CAS, these FLAT files are repackaged into 'CSV' (comma separated value) ASCII files which map 1-1 onto a SQLServer database table. The code to generate the CSV files is specially written, using the data model as a guide. There is potential for bugs here, and the CSV conversion code must be kept in sync with the data model (the FLAT files).

Load CAS: Once generated, the CSV files are loaded into the CAS, in general, they are incrementally added on an existing CAS, unless there are data model changes.

Build Indices: Once loaded, a Finish step is run to rebuild all indices for fast queries of objects.

Internal SDSS operations support

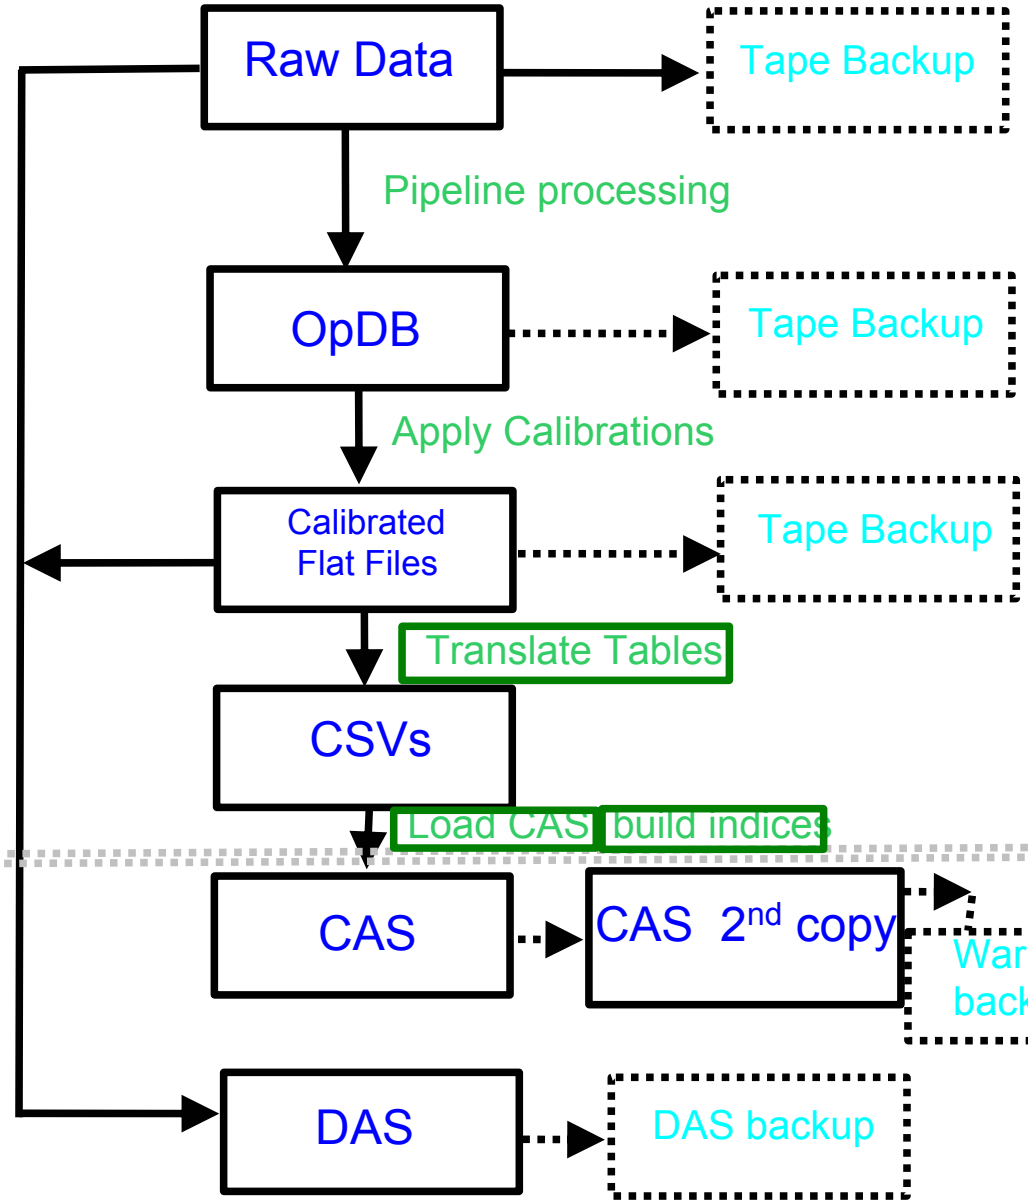

	DR1 2003	DR2 2004	DR5 (full survey) 2006+
	5TB	10TB	20TB (on tape)
	1TB	2TB	5TB per copy
	1TB	2TB	5TB per copy
	(1TB	2TB	5TB – ephemeral)
	1TB	2TB	5TB per copy
	[3 TB	6TB	15TB per copy]

External SDSS Data Distribution Support

Current status of hand-over to factory steps:

1. The 'Translate Tables' code (sqlFits2CSV) is done and working against the present data model. This code and routines have been checked into CVS (software code version control system), handed off to the 'FNAL data processing factory' and have been used successfully for DR1 and DR2 (and DR3 imaging). The shake out process required significant interaction with the developer and also required the addition of several 'data integrity checks'.

2. The 'Load CAS' code (for creating taskDBs) is checked into CVS and handed off to the 'FNAL data processing factory'. Incremental loading is supported for moving from one DR to the next; however, 'replacement' or 'changing' of individual imaging 'chunks' is not yet supported.

3. The 'build indices' code (also called 'the Finish step') is continuing to evolve. The version for DR2 differed substantially from that for DR1. We are aiming for stablization of this code for DR3. We are continuuing to interact closely with the developer for this piece of the code, and have not yet successfully run it 'independent of the developer' end-to-end.

Lessons Learned: (Terabyte flows of data are hard to keep laminar)

The size of the CSV files is similar to the size of the FLAT catalog files, about 1TB for DR1, 2TB for DR2, perhaps 5TB projected for the full survey by survey's end. Each copy of each version of the CAS itself is again similar in size.

These sizes don't fully include 'pixel image data' which makes DR1 DAS about 3TB total and DR2 about 6TB total.

Because of the difficulty in propogating changes from the data further upstream, The CAS works best as a 'read-only snapshot' of the state of the survey at a point in time about 1 year in the past.

This has generally proven (EDR,DR1) quite acceptable to the astronomical community to which we provide public access as part of our funding requirements.

This is generally not regarded as good enough for some of our astronomers within the collaboration who are eager to use the latest/largest/best-calibrated version of the data on a shorter timescale. In lieu of up-to-date CAS, we provide them with the 'DAS' on a much shorter timescale, although that data is more difficult to access (FLAT files, no indices, organized by chronological run number rather than position on the sky).

Internal SDSS operations support

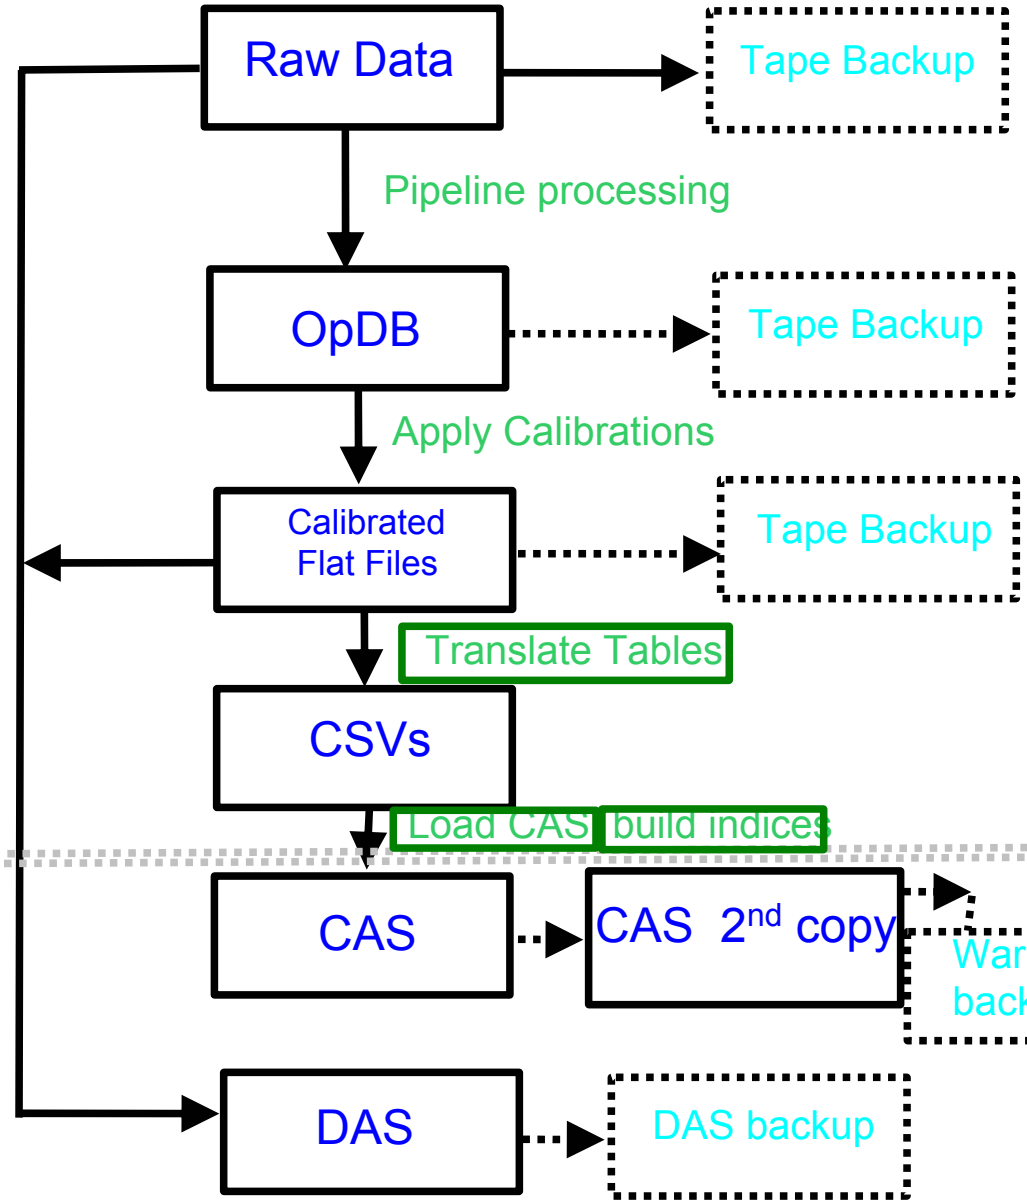

	DR1 2003	DR2 2004	DR5 (full survey) 2006+
	5TB	10TB	20TB (on tape)
OpDB	1TB	2TB	5TB per copy
Calibrated Flat Files	1TB	2TB	5TB per copy
CSVs	(1TB	2TB	5TB – ephemeral
CAS	1TB	2TB	5TB per copy
DAS	[3 TB	6TB	15TB per copy]

External SDSS Data Distribution Support

To keep a CAS database 'up to date' (say within one month of data collection) has been problematical because:

1. A trickle of changing data model and upstream processing changes, propogates to a torrent downstream.

   There is a constant stream of 'problem reports' regard one or another aspect of the data, either in the DAS or the CAS. These are generally acknowledged fairly quickly, and a fix is generally straight-forward FOR FUTURE processings or FUTURE releases of the data, but to propogate a fix downstream into the 'current' active release is often very difficult and time-consuming.

2. Limited disk space

   Recall we have to keep usually 6 times the size of the original: raw data,OPDB,FLAT files, CSVs, CAS, CAS backup spinning). Thus a 2TB DR2 actually requires about 12 TB spinning disk. Just copying 1TB around usually requires about 1 day at current network/disk I/O speeds.

3. Regeneration of 'full-indices' (spectro photo match, incremental loading with full indicies) takes of order 2 weeks time within the CAS. And the definition of what's indexed is still changing somewhat.

Example of how syncing a CAS for a Data Release actually goes....

Sample of 'Patches' made to DR2 CAS: (Scheduled for Release to public Mar 15, 2004)

2003 Dec 1: Preliminary DR2 CAS announced to archive testing group

2003 Dec 2: Report of missing galaxies in portion of DR2 footprint
Problem traced to bug upstream in 'resolve the sky out of the OPDB code'
Code fixed upstream, patch generated (~10Million CSV records) applied

2003 Dec 15: Preliminary DR2 CAS announced to collaboration

2004 Jan 25: Reports of missing target, tiling matches

2004 Feb 5: Reports of improperly masked regions of bad data

2004 Feb 9: Reports of missing USNO matches

2004 Feb 13: Three more patches applied, subset of Finish step rerun 'by hand'

2004 Feb 18: Report of missing photometric information on 'bonus' plates beyond survey boundary

2004 Feb 20: Patch applied

2004 Feb 25: 'Final' DR2 CAS announced to collaboration

2004 Feb 26: Report that patch applied on Feb 20th didn't solve whole problem...

Future plans:

Push forward with DR3 (Oct 2004), DR4 and beyond.
Triage: 'bugs', 'features', 'documentation issues'.

Continue to minimize data model changes, coordinate data model changes in sync with data release fixed points.

Try and document and 'codify' all knowledge learned from DR2 patching and loading so that it's carried forward to DR3 and beyond.

Look at options for 'deleting and replacing' a subset of data in a database. (more formalized patch process).

Look at options for 'updating the calibrations' without reloading the whole database.

Further future:

Explore the possibility of a limited 'Runs database' for the collaboration which contains 'nearly all' the SDSS imaging and spectra on a shorter timescale.